

Powering Hidden Markov Model by Neural Network based Generative Models

Dong Liu[†] and Antoine Honoré^{*†} and Saikat Chatterjee[†] and Lars K. Rasmussen^{† 1}

Abstract. Hidden Markov model (HMM) has been successfully used for sequential data modeling problems. In this work, we propose to power the modeling capacity of HMM by bringing in neural network based generative models. The proposed model is termed as GenHMM. In the proposed GenHMM, each HMM hidden state is associated with a neural network based generative model that has tractability of exact likelihood and provides efficient likelihood computation. A generative model in GenHMM consists of a mixture of generators that are realized by flow models. A learning algorithm for GenHMM is proposed in expectation-maximization framework. The convergence of the learning GenHMM is analyzed. We demonstrate the efficiency of GenHMM by classification tasks on practical sequential data.

1 Introduction

Sequential data modeling is a challenging topic in pattern recognition and machine learning. For many applications, the assumption of independent and identically distributed (i.i.d.) data points is too strong to model data properly. Hidden Markov model (HMM) is a classic way to model sequential data without the i.i.d. assumption. HMM has been widely used in different practical problems, including applications in reinforcement learning [7, 19], natural language modeling [15, 12], biological sequence analysis such as proteins [1] and DNA [24], etc.

A HMM is a statistical representation of sequential data generating process. Each state of a HMM is associated with a probabilistic model. The probabilistic model is used to represent the relationship between a state of HMM and sequential data input. The typical way is to use a Gaussian mixture model (GMM) per state of HMM [2], where GMMs are used to connect states of HMM to sequential data input. GMM based HMM (GMM-HMM) has become a standard model for sequential data modeling, and been employed widely for practical applications, especially in speech recognition [10, 5].

Given the success of GMM-HMM, it is not efficient for modeling data in nonlinear manifold. Research attempts at training HMM with neural networks have been made to boost the modeling capacity of HMM. A successful work of this track has brought deep neural network (DNN) that is defined by restrictive Boltzmann machines (RBMs) [14] into HMM based models [13, 20, 23]. RBM based HMM is trained with a hierarchical scheme consisting of multiple steps of unsupervised learning, formatting of a classification network and then supervised learning. The hierarchical procedure comes from the empirical expertise in this domain. To be more specific, the hier-

archical learning scheme of RBM/DNN based HMM consists of: i) RBMs are trained one after the other in unsupervised fashion, and are stacked together as one deep neural network model, ii) then a final softmax layer is added to the stack of RBMs to represent the probability of a HMM state given a data input, iii) a discriminative training is performed for the final tuning of the model at the final stage.

Another track of related work is hybrid method of temporal neural network models and HMM. In [21, 4, 18], a long short-term memory (LSTM) model/recurrent neural network (RNN) is combined with HMM as hybrid. A hierarchical training is carried out by: i) training a HMM first, ii) then doing modified training of LSTM using trained HMM. This hierarchical training procedure is motivated by the intuition of using LSTM or RNN to fill in the gap where HMM can not learn.

The above works help improve modeling capacity of HMM based models by bringing in neural networks. A softmax layer is usually used to represent probability whenever a conditional distribution is needed. These hierarchical schemes are built based on intuition of domain knowledge. Training of these hierarchical models usually requires expertise in specific areas to be able to proceed with the hierarchical procedure of training and application usage.

In this work, we propose a generative model based HMM, termed as GenHMM. Specifically, a generative model in our GenHMM is generator-mixed, where a generator is realized by a neural network to help the model gain high modeling capacity. Our proposed model, GenHMM,

- has high modeling capacity of sequential data, due to the neural network based generators;
- is easy to train. Training of GenHMM employs expectation maximization (EM) framework. Therefore, training a GenHMM is as easy as training a GMM-HMM model, while configuration of GenHMM is flexible;
- is able to compute loglikelihood exactly and efficiently.

Instead of using softmax for probability representation, our GenHMM has tractability of exact loglikelihood of given sequential data, which is based on the change of variable formula. To make the loglikelihood computation efficient, neural network based generators of GenHMM are realized as flow models.

Our contributions in the paper are as follows.

- Proposing a neural network based HMM for sequential data modeling, i.e. GenHMM. GenHMM has the tractability of exact likelihood.
- Designing practical algorithm for training GenHMM under EM framework. Stochastic gradient search in batch fashion is embedded in this algorithm.

[†] KTH Royal Institute of Technology, Stockholm, Sweden

*Karolinska Institute, Stockholm, Sweden

E-mail: {doli, honore, sach, lkra}@kth.se

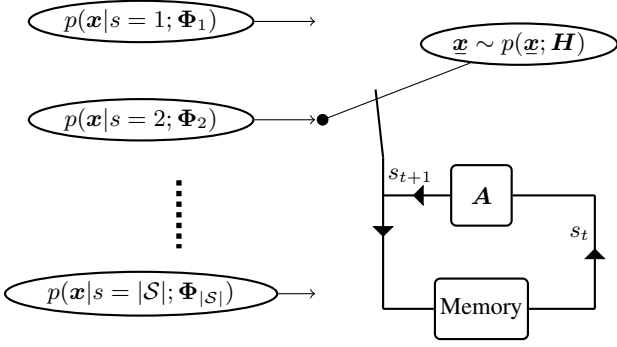


Figure 1: HMM model illustration.

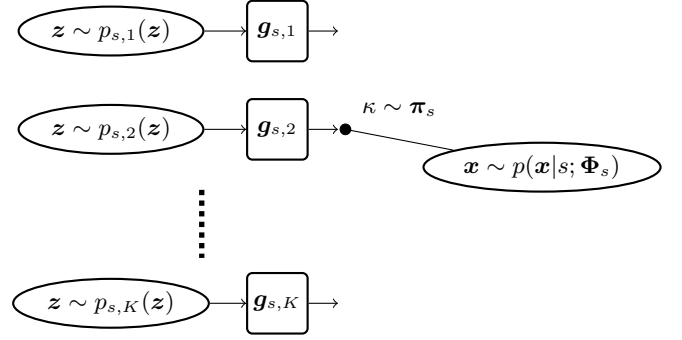


Figure 2: Source of state s in GenHMM.

- Giving convergence analysis for GenHMM under the proposed learning algorithm.
- Verifying the proposed model on practical sequential data.

2 Generator-mixed HMM (GenHMM)

Our framework is a HMM. A HMM \mathbf{H} defined in a hypothesis space \mathcal{H} , i.e. $\mathbf{H} \in \mathcal{H}$, is capable to model time-span signal $\mathbf{x} = [\mathbf{x}_1, \dots, \mathbf{x}_T]^\top$, where $\mathbf{x}_t \in \mathbb{R}^N$ is the N -dimensional signal at time t , $[\cdot]^\top$ denotes transpose, and T denotes the time length². We define the hypothesis set of HMM as $\mathcal{H} := \{\mathcal{S}, \mathbf{q}, \mathbf{A}, p(\mathbf{x}|\mathbf{s}; \Phi_s)\}$, where

- \mathcal{S} is the set of hidden states of \mathbf{H} .
- $\mathbf{q} = [q_1, q_2, \dots, q_{|\mathcal{S}|}]^\top$ is the initial state distribution of \mathbf{H} with $|\mathcal{S}|$ as cardinality of \mathcal{S} . For $i \in \mathcal{S}$, $q_i = p(s_1 = i; \mathbf{H})$. We use s_t to denote the state s at time t .
- \mathbf{A} matrix of size $|\mathcal{S}| \times |\mathcal{S}|$ is the transition matrix of states in \mathbf{H} . That is, $\forall i, j \in \mathcal{S}, \mathbf{A}_{i,j} = p(s_{t+1} = j | s_t = i; \mathbf{H})$.
- For a given hidden state s , the density function of the observable signal is $p(\mathbf{x}|s; \Phi_s)$, where Φ_s is the parameter set that defines this probabilistic model. Denote $\Phi = \{\Phi_s | s \in \mathcal{S}\}$.

Using HMM for signal representation is illustrated in Figure 1. The model assumption is that different instant signal of \mathbf{x} is generated by a different signal source associated with a hidden state of HMM. In the framework of HMM, at each time instance t , signal \mathbf{x}_t is assumed to be generated by a distribution with density function $p(\mathbf{x}_t | s_t; \Phi_{s_t})$, and s_t is decided by the hidden markov process. Putting these together gives us the probabilistic model $p(\mathbf{x}; \mathbf{H})$.

2.1 Generative Model of GenHMM

In this section, we introduce the neural network based state probabilistic model of our GenHMM. Recall that $\mathbf{x} \in \mathbb{R}^N$. Subscript is omitted when it does not cause ambiguity. The probabilistic model of GenHMM for each hidden state is a mixture of K neural network based generators, where K is a positive integer. The probabilistic model of a state $s \in \mathcal{S}$ is then given by

$$p(\mathbf{x}|s; \Phi_s) = \sum_{\kappa=1}^K \pi_{s,\kappa} p(\mathbf{x}|s, \kappa; \theta_{s,\kappa}), \quad (1)$$

where κ is a random variable following a categorical distribution, with probability $\pi_{s,\kappa} = p(\kappa|s; \mathbf{H})$. Naturally $\sum_{\kappa=1}^K \pi_{s,\kappa} = 1$. Denote $\boldsymbol{\pi}_s = [\pi_{s,1}, \pi_{s,2}, \dots, \pi_{s,K}]^\top$. In (1), $p(\mathbf{x}|s, \kappa; \theta_{s,\kappa})$ is defined

as induced distribution by a generator $\mathbf{g}_{s,\kappa} : \mathbb{R}^N \rightarrow \mathbb{R}^N$, such that $\mathbf{x} = \mathbf{g}_{s,\kappa}(\mathbf{z})$, where \mathbf{z} is a latent variable following a distribution with density function $p_{s,\kappa}(\mathbf{z})$. Generator $\mathbf{g}_{s,\kappa}$ is parameterized by $\theta_{s,\kappa}$. Let us denote the collection of the parameter sets of generators for state s as $\theta_s = \{\theta_{s,\kappa} | \kappa = 1, 2, \dots, K\}$. Assuming $\mathbf{g}_{s,\kappa}$ is invertible, by change of variable, we have

$$p(\mathbf{x}|s, \kappa; \theta_{s,\kappa}) = p_{s,\kappa}(\mathbf{z}) \left| \det \left(\frac{\partial \mathbf{g}_{s,\kappa}(\mathbf{z})}{\partial \mathbf{z}} \right) \right|^{-1}. \quad (2)$$

The signal flow of the probability distribution for a state s of GenHMM is shown in Figure 2, in which the generator identity is up to the random variable κ .

2.2 Learning in EM framework

Assume the sequential signal \mathbf{x} follows an unknown distribution $p(\mathbf{x})$. We would like to use GenHMM to model this distribution. Alternatively, we are looking for the answer to the question

$$\min_{\mathbf{H} \in \mathcal{H}} KL(p(\mathbf{x}) \| p(\mathbf{x}; \mathbf{H})), \quad (3)$$

where $KL(\cdot \| \cdot)$ denotes the Kullback-Leibler divergence. For practical consideration, we only have access to the samples of $p(\mathbf{x})$, i.e. the dataset of this distribution. For the given dataset, we denote its empirical distribution by $\hat{p}(\mathbf{x}) = \frac{1}{R} \sum_{r=1}^R \delta_{\mathbf{x}^r}(\mathbf{x})$, where R denotes the total number of sequential samples and superscript $(\cdot)^r$ denotes the index of r -th sequential signal. The KL divergence minimization problem can be reduced to a likelihood maximization problem

$$\operatorname{argmax}_{\mathbf{H} \in \mathcal{H}} \frac{1}{R} \sum_{r=1}^R \log p(\mathbf{x}^r; \mathbf{H}). \quad (4)$$

For the likelihood maximization, the first problem that we need to address is to deal with the hidden sequential variables of model \mathbf{H} , namely $\mathbf{s} = [s_1, s_2, \dots, s_T]^\top$ and $\boldsymbol{\kappa} = [\kappa_1, \kappa_2, \dots, \kappa_T]^\top$. For a sequential observable variable \mathbf{x} , \mathbf{s} is the hidden state sequence corresponding to \mathbf{x} , and $\boldsymbol{\kappa}$ is the hidden variable sequence representing the generator identity sequence that actually generates \mathbf{x} .

Since directly maximizing likelihood is not an option for our problem in (4), we address this problem in expectation maximization (EM) framework. This divides our problem into two iterative steps: i) using the joint posterior of hidden variable sequences \mathbf{s} and $\boldsymbol{\kappa}$ to obtain an ‘‘expected likelihood’’ of the observable variable sequence \mathbf{x} , i.e. the E-step; ii) maximizing the expected likelihood with regard to (w.r.t.) the model \mathbf{H} , i.e. the M-step. Assume model \mathbf{H} is at a configuration of \mathbf{H}^{old} , we formulate these two steps as follows.

² The length for sequential data varies.

- E-step: the expected likelihood function

$$\mathcal{Q}(\mathbf{H}; \mathbf{H}^{\text{old}}) = \mathbb{E}_{\hat{p}(\mathbf{x}), p(\mathbf{s}, \kappa | \mathbf{x}; \mathbf{H}^{\text{old}})} [\log p(\mathbf{x}, \mathbf{s}, \kappa; \mathbf{H})], \quad (5)$$

where $\mathbb{E}_{\hat{p}(\mathbf{x}), p(\mathbf{s}, \kappa | \mathbf{x}; \mathbf{H}^{\text{old}})} [\cdot]$ denotes the expectation operator by distribution $\hat{p}(\mathbf{x})$ and $p(\mathbf{s}, \kappa | \mathbf{x}; \mathbf{H}^{\text{old}})$.

- M-step: the maximization step

$$\max_{\mathbf{H}} \mathcal{Q}(\mathbf{H}; \mathbf{H}^{\text{old}}). \quad (6)$$

The problem (6) can be reformulated as

$$\begin{aligned} & \max_{\mathbf{H}} \mathcal{Q}(\mathbf{H}; \mathbf{H}^{\text{old}}) \\ = & \max_{\mathbf{q}} \mathcal{Q}(\mathbf{q}; \mathbf{H}^{\text{old}}) + \max_{\mathbf{A}} \mathcal{Q}(\mathbf{A}; \mathbf{H}^{\text{old}}) + \max_{\Phi} \mathcal{Q}(\Phi; \mathbf{H}^{\text{old}}), \end{aligned} \quad (7)$$

where the decomposed optimization problems are

$$\mathcal{Q}(\mathbf{q}; \mathbf{H}^{\text{old}}) = \mathbb{E}_{\hat{p}(\mathbf{x}), p(\mathbf{s} | \mathbf{x}; \mathbf{H}^{\text{old}})} [\log p(s_1; \mathbf{H})], \quad (8)$$

$$\mathcal{Q}(\mathbf{A}; \mathbf{H}^{\text{old}}) = \mathbb{E}_{\hat{p}(\mathbf{x}), p(\mathbf{s} | \mathbf{x}; \mathbf{H}^{\text{old}})} \left[\sum_{t=1}^{T-1} \log p(s_{t+1} | s_t; \mathbf{H}) \right], \quad (9)$$

$$\mathcal{Q}(\Phi; \mathbf{H}^{\text{old}}) = \mathbb{E}_{\hat{p}(\mathbf{x}), p(\mathbf{s}, \kappa | \mathbf{x}; \mathbf{H}^{\text{old}})} [\log p(\mathbf{x}, \kappa | \mathbf{s}; \mathbf{H})]. \quad (10)$$

We can see that the solution of \mathbf{H} depends on the posterior probability $p(\mathbf{s} | \mathbf{x}; \mathbf{H})$. Though the evaluation of posterior according to Bayesian theorem is straightforward, the computation complexity of $p(\mathbf{s} | \mathbf{x}; \mathbf{H})$ grows exponentially with the length of \mathbf{s} . Therefore, we employ forward-backward algorithm [3] to do the posterior computation efficiently. As we would detail in the next section, what are needed to formulate the problem, are actually the $p(\mathbf{s} | \mathbf{x}; \mathbf{H})$ and $p(s, \kappa | \mathbf{x}; \mathbf{H})$. For the joint posterior $p(\mathbf{s}, \kappa | \mathbf{x}; \mathbf{H})$, it can be computed by the Bayesian rule when posterior of hidden state is available.

With such a solution framework ready for GenHMM, there are still remaining problems to address before it can be employed for practical usage, including

- how to realize GenHMM by neural network based generators such that likelihood of their induced distributions can be computed explicitly and exactly?
- how to train GenHMM to solve problem in (4) using practical algorithm?
- would the training of GenHMM converge?

We tackle these problems in the following section.

3 Solution for GenHMM

In this section, we detail the solution for realizing and learning GenHMM. The convergence of GenHMM is also discussed in this section.

3.1 Realizing $g_{s, \kappa}$ by a Flow Model

Each generator $g_{s, \kappa}$ is realized as a feed-forward neural network. We define $g_{s, \kappa}$ as a L -layer neural network and formulate its mapping by layer-wise concatenation: $g_{s, \kappa} = g_{s, \kappa}^{[L]} \circ g_{s, \kappa}^{[L-1]} \circ \dots \circ g_{s, \kappa}^{[1]}$, where superscript $[l]$ denotes the layer index and \circ denotes mapping concatenation. Assume $g_{s, \kappa}$ is invertible and denote its inverse mapping as $f_{s, \kappa} = g_{s, \kappa}^{-1}$. For a latent variable \mathbf{z} with density function $p_{s, \kappa}(\mathbf{z})$, the generated signal \mathbf{x} follows an induced distribution with

density function (2). We illustrate the signal flow between latent variable \mathbf{z} and observable variable \mathbf{x} as

$$\mathbf{z} = \mathbf{h}_0 \begin{array}{c} \xrightarrow{g_{s, \kappa}^{[1]}} \\ \xleftarrow{f_{s, \kappa}^{[1]}} \end{array} \mathbf{h}_1 \begin{array}{c} \xrightarrow{g_{s, \kappa}^{[2]}} \\ \xleftarrow{f_{s, \kappa}^{[2]}} \end{array} \dots \dots \dots \begin{array}{c} \xrightarrow{g_{s, \kappa}^{[L]}} \\ \xleftarrow{f_{s, \kappa}^{[L]}} \end{array} \mathbf{h}_L = \mathbf{x}$$

where $f_{s, \kappa}^{[l]}$ is the l -th layer of $f_{s, \kappa}$. We have $\mathbf{z} = f_{s, \kappa}(\mathbf{x})$. If every layer of $g_{s, \kappa}$ is invertible, the full feed-forward neural network is invertible. Flow model, proposed in [9] as an image generating model, is such an invertible feed-forward layer-wise neural network. It is further improved in subsequential works [8, 17] for high-fidelity and high-resolution image generating and representation. As shown in (2), the challenge lies at the computation of Jacobian determinant. Another track of flow models uses a continuous-depth models instead. The variable change is defined by an ordinary differential equation implemented by a neural network [6, 11], where the key becomes to solve the ODE problem. We use the layer-wise flow model to model the variable change in (2) in which the efficient Jacobian computation is available.

For a flow model, let us assume that the feature \mathbf{h}_l at the l 'th layer has two subparts as $\mathbf{h}_l = [\mathbf{h}_{l,a}^T, \mathbf{h}_{l,b}^T]^T$. The efficient invertible mapping of flow model comes from following forward and inverse relations between $(l-1)$ 'th and l 'th layers

$$\begin{aligned} \mathbf{h}_l &= \begin{bmatrix} \mathbf{h}_{l,a} \\ \mathbf{h}_{l,b} \end{bmatrix} = \begin{bmatrix} \mathbf{h}_{l-1,a} \\ (\mathbf{h}_{l-1,b} - \mathbf{m}_b(\mathbf{h}_{l-1,a})) \odot \mathbf{m}_a(\mathbf{h}_{l-1,a}) \end{bmatrix}, \\ \mathbf{h}_{l-1} &= \begin{bmatrix} \mathbf{h}_{l-1,a} \\ \mathbf{h}_{l-1,b} \end{bmatrix} = \begin{bmatrix} \mathbf{h}_{l,a} \\ \mathbf{m}_a(\mathbf{h}_{l,a}) \odot \mathbf{h}_{l,b} + \mathbf{m}_b(\mathbf{h}_{l,a}) \end{bmatrix}, \end{aligned} \quad (11)$$

where \odot denotes element-wise product, \oslash denotes element-wise division, and $\mathbf{m}_a(\cdot)$, $\mathbf{m}_b(\cdot)$ can be complex non-linear mappings (implemented by neural networks). For the flow model, the determinant of Jacobian matrix is

$$\det(\nabla f_{s, \kappa}) = \prod_{l=1}^L \det(\nabla f_{s, \kappa}^{[l]}), \quad (12)$$

where $\nabla f_{s, \kappa}^{[l]}$ is the Jacobian of the mapping from the l -th layer to the $(l-1)$ -th layer, i.e., the inverse transformation. We compute the determinant of the Jacobian matrix as

$$\begin{aligned} \det(\nabla f_{s, \kappa}^{[l]}) &= \det \begin{bmatrix} \frac{\partial \mathbf{h}_{l-1}}{\partial \mathbf{h}_l} \end{bmatrix} \\ &= \det \begin{bmatrix} \mathbf{I}_a & \mathbf{0} \\ \frac{\partial \mathbf{h}_{l-1,b}}{\partial \mathbf{h}_{l,a}} & \text{diag}(\mathbf{m}_a(\mathbf{h}_{l,a})) \end{bmatrix} \\ &= \det(\text{diag}(\mathbf{m}_a(\mathbf{h}_{l,a}))), \end{aligned} \quad (13)$$

where \mathbf{I}_a is identity matrix and $\text{diag}(\cdot)$ returns a square matrix with the elements of \cdot on the main diagonal.

(11) describes a *coupling* layer in a flow model. A flow model is basically a stack of multiple coupling layers. But the issue of direct concatenation of multiple such coupling mappings is partial identity mapping of the whole model. This issue can be addressed by alternating hidden signal order after each coupling layer.

3.2 Learning of GenHMM

In this subsection, we address the problem of learning GenHMM.

3.2.1 Generative Model Learning

The generative model learning is actually to solve the problem in (10), which can be further divided into two subproblems: i) generator learning; ii) mixture weights of generators learning. Let us define notations: $\mathbf{\Pi} = \{\pi_s | s \in \mathcal{S}\}$, $\mathbf{\Theta} = \{\theta_s | s \in \mathcal{S}\}$. Then the problem in (10) becomes

$$\max_{\mathbf{\Phi}} \mathcal{Q}(\mathbf{\Phi}; \mathbf{H}^{\text{old}}) = \max_{\mathbf{\Pi}} \mathcal{Q}(\mathbf{\Pi}; \mathbf{H}^{\text{old}}) + \max_{\mathbf{\Theta}} \mathcal{Q}(\mathbf{\Theta}; \mathbf{H}^{\text{old}}), \quad (14)$$

where

$$\mathcal{Q}(\mathbf{\Pi}; \mathbf{H}^{\text{old}}) = \mathbb{E}_{\hat{p}(\mathbf{x}), p(\mathbf{s}, \mathbf{\kappa} | \mathbf{x}; \mathbf{H}^{\text{old}})} [\log p(\mathbf{\kappa} | \mathbf{s}; \mathbf{H})], \quad (15)$$

$$\mathcal{Q}(\mathbf{\Theta}; \mathbf{H}^{\text{old}}) = \mathbb{E}_{\hat{p}(\mathbf{x}), p(\mathbf{s}, \mathbf{\kappa} | \mathbf{x}; \mathbf{H}^{\text{old}})} [\log p(\mathbf{x} | \mathbf{s}, \mathbf{\kappa}; \mathbf{H})]. \quad (16)$$

We firstly address the generator learning problem, i.e. $\max_{\mathbf{\Theta}} \mathcal{Q}(\mathbf{\Theta}; \mathbf{H}^{\text{old}})$. This is boiled down to maximize the cost function of neural networks that can be formulated as

$$\begin{aligned} & \mathcal{Q}(\mathbf{\Theta}; \mathbf{H}^{\text{old}}) \\ &= \frac{1}{R} \sum_{r=1}^R \sum_{\mathbf{s}^r} \sum_{\mathbf{\kappa}^r} p(\mathbf{s}^r, \mathbf{\kappa}^r | \mathbf{x}^r; \mathbf{H}^{\text{old}}) \sum_{t=1}^{T^r} \log p(\mathbf{x}_t^r | s_t^r, \kappa_t^r; \mathbf{H}) \\ &= \frac{1}{R} \sum_{r=1}^R \sum_{t=1}^{T^r} \sum_{s_t^r=1}^{|\mathcal{S}|} \sum_{\kappa_t^r=1}^K p(s_t^r | \mathbf{x}^r; \mathbf{H}^{\text{old}}) p(\kappa_t^r | s_t^r, \mathbf{x}^r; \mathbf{H}^{\text{old}}) \\ & \quad \log p(\mathbf{x}_t^r | s_t^r, \kappa_t^r; \mathbf{H}), \end{aligned} \quad (17)$$

where T^r is the length of the r -th sequential data. In (17), the state posterior $p(s_t | \mathbf{x}, \mathbf{H}^{\text{old}})$ is computed by forward-backward algorithm. The posterior of κ is

$$\begin{aligned} p(\kappa | s, \mathbf{x}; \mathbf{H}^{\text{old}}) &= \frac{p(\kappa, \mathbf{x} | s; \mathbf{H}^{\text{old}})}{p(\mathbf{x} | s; \mathbf{H}^{\text{old}})} \\ &= \frac{\pi_{s, \kappa}^{\text{old}} p(\mathbf{x} | s, \kappa; \mathbf{H}^{\text{old}})}{\sum_{\kappa=1}^K \pi_{s, \kappa}^{\text{old}} p(\mathbf{x} | s, \kappa; \mathbf{H}^{\text{old}})}, \end{aligned} \quad (18)$$

where the last equation is due to the fact that \mathbf{x}_t among sequence \mathbf{x} only depends on s_t, κ_t .

By substituting (2) and (12) into (17), we have cost function for neural networks as

$$\begin{aligned} & \mathcal{Q}(\mathbf{\Theta}; \mathbf{H}^{\text{old}}) \\ &= \frac{1}{R} \sum_{r=1}^R \sum_{t=1}^{T^r} \sum_{s_t^r=1}^{|\mathcal{S}|} \sum_{\kappa_t^r=1}^K p(s_t^r | \mathbf{x}^r; \mathbf{H}^{\text{old}}) p(\kappa_t^r | s_t^r, \mathbf{x}^r; \mathbf{H}^{\text{old}}) \\ & \quad \left[\log p_{s_t^r, \kappa_t^r}(\mathbf{f}_{s_t^r, \kappa_t^r}(\mathbf{x}_t^r)) + \sum_{l=1}^L \log |\det(\nabla \mathbf{f}_{s, \kappa}^{[l]})| \right]. \end{aligned} \quad (19)$$

The generators of GenHMM simply use standard Gaussian distribution for latent variables $\mathbf{z} \sim p_{s, \kappa}(\mathbf{z})$. Since training dataset can be too large to do whole-dataset iterations, batch-size stochastic gradient decent can be used to maximize $\mathcal{Q}(\mathbf{\Theta}; \mathbf{H}^{\text{old}})$ w.r.t. parameters of generators.

In what follows we address the problem $\max_{\mathbf{\Pi}} \mathcal{Q}(\mathbf{\Pi}; \mathbf{H}^{\text{old}})$ in our generative model learning. The conditional distribution of hidden variable κ , $\pi_{s, \kappa} = p(\kappa | s; \mathbf{H})$, is obtained by solving the following problem

$$\begin{aligned} \pi_{s, \kappa} &= \operatorname{argmax}_{\pi_{s, \kappa}} \mathcal{Q}(\mathbf{\Pi}; \mathbf{H}^{\text{old}}) \quad (20) \\ \text{s.t.} \quad & \sum_{\kappa=1}^K \pi_{s, \kappa} = 1, \forall s = 1, 2, \dots, |\mathcal{S}|. \end{aligned}$$

To solve problem (20), we formulate its Lagrange function as

$$\mathcal{L} = \mathcal{Q}(\mathbf{\Pi}; \mathbf{H}^{\text{old}}) + \sum_{s=1}^{|\mathcal{S}|} \lambda_s \left(1 - \sum_{\kappa=1}^K \pi_{s, \kappa} \right). \quad (21)$$

Solving $\frac{\partial \mathcal{L}}{\partial \pi_{s, \kappa}} = 0$ gives

$$\pi_{s, \kappa} = \frac{1}{\lambda_s} \sum_{r=1}^R \sum_{t=1}^{T^r} p(s_t^r = s, \kappa_t^r = \kappa | \mathbf{x}^r; \mathbf{H}^{\text{old}}). \quad (22)$$

With condition $\sum_{\kappa=1}^K \pi_{s, \kappa} = 1, \forall s = 1, 2, \dots, |\mathcal{S}|$, we have

$$\lambda_s = \sum_{\kappa=1}^K \sum_{r=1}^R \sum_{t=1}^{T^r} p(s_t^r = s, \kappa_t^r = \kappa | \mathbf{x}^r; \mathbf{H}^{\text{old}}). \quad (23)$$

Then the solution to (20) is

$$\pi_{s, \kappa} = \frac{\sum_{r=1}^R \sum_{t=1}^{T^r} p(s_t^r = s, \kappa_t^r = \kappa | \mathbf{x}^r; \mathbf{H}^{\text{old}})}{\sum_{\kappa=1}^K \sum_{r=1}^R \sum_{t=1}^{T^r} p(s_t^r = s, \kappa_t^r = \kappa | \mathbf{x}^r; \mathbf{H}^{\text{old}})}, \quad (24)$$

where

$$p(s, \kappa | \mathbf{x}; \mathbf{H}^{\text{old}}) = p(s | \mathbf{x}; \mathbf{H}^{\text{old}}) p(\kappa | s, \mathbf{x}; \mathbf{H}^{\text{old}}). \quad (25)$$

Here $p(s | \mathbf{x}; \mathbf{H}^{\text{old}})$ can be computed by forward-backward algorithm, while $p(\kappa | s, \mathbf{x}; \mathbf{H}^{\text{old}})$ is given by (18).

With the generative model learning obtained, it remains to solve the initial distribution update and transition matrix update of HMM in GenHMM, i.e. the problem (8) and (9). These two problems are basically two constrained optimization problems. The solutions to them are available in literature [3]. But to keep learning algorithm for GenHMM complete, we give the update rules for \mathbf{q} and \mathbf{A} as follows.

3.2.2 Initial Probability Update

The problem in (8) can be reformulated as

$$\begin{aligned} & \mathcal{Q}(\mathbf{q}; \mathbf{H}^{\text{old}}) \\ &= \frac{1}{R} \sum_{r=1}^R \sum_{\mathbf{s}^r} p(\mathbf{s}^r | \mathbf{x}^r; \mathbf{H}^{\text{old}}) \log p(s_1^r; \mathbf{H}) \\ &= \frac{1}{R} \sum_{r=1}^R \sum_{s_1^r=1}^{|\mathcal{S}|} \sum_{s_2^r=1}^{|\mathcal{S}|} \dots \sum_{s_{T^r}^r=1}^{|\mathcal{S}|} p(s_1^r, s_2^r, \dots, s_{T^r}^r | \mathbf{x}^r; \mathbf{H}^{\text{old}}) \log p(s_1^r) \\ &= \frac{1}{R} \sum_{r=1}^R \sum_{s_1^r=1}^{|\mathcal{S}|} p(s_1^r | \mathbf{x}^r; \mathbf{H}^{\text{old}}) \log p(s_1^r; \mathbf{H}). \end{aligned} \quad (26)$$

$p(s_1^r; \mathbf{H})$ is the probability of initial state of GenHMM for r -th sequential sample. Actually $q_i = p(s_1 = i; \mathbf{H}), i = 1, 2, \dots, |\mathcal{S}|$. Solution to the problem

$$\mathbf{q} = \operatorname{argmax}_{\mathbf{q}} \mathcal{Q}(\mathbf{q}; \mathbf{H}^{\text{old}}), \text{ s.t. } \sum_{i=1}^{|\mathcal{S}|} q_i = 1, q_i \geq 0, \forall i. \quad (27)$$

is

$$q_i = \frac{1}{R} \sum_{r=1}^R p(s_1^r = i | \mathbf{x}^r; \mathbf{H}^{\text{old}}), \forall i = 1, 2, \dots, |\mathcal{S}|. \quad (28)$$

3.2.3 Transition Probability Update

The problem (9) can be reformulated as

$$\begin{aligned} & \mathcal{Q}(\mathbf{A}; \mathbf{H}^{\text{old}}) \\ &= \sum_{r=1}^R \sum_{\mathbf{s}^r} p(\mathbf{s}^r | \mathbf{x}^r; \mathbf{H}^{\text{old}}) \sum_{t=1}^{T^r-1} \log p(s_{t+1}^r | s_t^r; \mathbf{H}) \\ &= \sum_{r=1}^R \sum_{t=1}^{T^r-1} \sum_{s_t^r=1}^{|\mathcal{S}|} \sum_{s_{t+1}^r=1}^{|\mathcal{S}|} p(s_t^r, s_{t+1}^r | \mathbf{x}^r; \mathbf{H}^{\text{old}}) \log p(s_{t+1}^r | s_t^r; \mathbf{H}). \end{aligned} \quad (29)$$

Since $\mathbf{A}_{i,j} = p(s_{t+1}^r = j | s_t^r = i; \mathbf{H})$ is the element of transition matrix \mathbf{A} , the solution to the problem

$$\begin{aligned} & \mathbf{A} = \underset{\mathbf{A}}{\text{argmax}} \mathcal{Q}(\mathbf{A}; \mathbf{H}^{\text{old}}) \\ & \text{s.t. } \mathbf{A} \cdot \mathbf{1} = \mathbf{1}, \mathbf{A}_{i,j} \geq 0 \forall i, j, \end{aligned} \quad (30)$$

is

$$\mathbf{A}_{i,j} = \frac{\bar{\xi}_{i,j}}{\sum_{k=1}^{|\mathcal{S}|} \bar{\xi}_{i,k}}, \quad (31)$$

where

$$\bar{\xi}_{i,j} = \sum_{r=1}^R \sum_{t=1}^{T^r-1} p(s_t^r = i, s_{t+1}^r = j | \mathbf{x}^r; \mathbf{H}^{\text{old}}). \quad (32)$$

3.3 On Convergence of GenHMM

In pursuit of representing a dataset by GenHMM, we are interested if the learning solution discussed in subsection 3.2 would converge. The properties on GenHMM's convergence are analyzed as follows.

Proposition 1. Assume that parameter $\Theta = \{\theta_{s,\kappa} | s \in \mathcal{S}, \kappa = 1, 2, \dots, K\}$ is in a compact set, $\mathbf{f}_{s,\kappa}$ and $\nabla \mathbf{f}_{s,\kappa}$ are continuous w.r.t. $\theta_{s,\kappa}$ in GenHMM. Then GenHMM converges.

Proof. We begin with the comparison of loglikelihood evaluated under \mathbf{H}^{new} and \mathbf{H}^{old} . The loglikelihood of dataset given by $\hat{p}(\mathbf{x})$ can be reformulated as

$$\begin{aligned} & \mathbb{E}_{\hat{p}(\mathbf{x})} [\log p(\mathbf{x}; \mathbf{H}^{\text{new}})] \\ &= \mathbb{E}_{\hat{p}(\mathbf{x}), p(\mathbf{s}, \kappa | \mathbf{x}; \mathbf{H}^{\text{old}})} \left[\log \frac{p(\mathbf{x}, \mathbf{s}, \kappa; \mathbf{H}^{\text{new}})}{p(\mathbf{s}, \kappa | \mathbf{x}; \mathbf{H}^{\text{old}})} \right] \\ &+ \mathbb{E}_{\hat{p}(\mathbf{x})} \left[KL(p(\mathbf{s}, \kappa | \mathbf{x}; \mathbf{H}^{\text{old}}) || p(\mathbf{s}, \kappa | \mathbf{x}; \mathbf{H}^{\text{new}})) \right], \end{aligned}$$

where the first term on the right hand side of the above inequality can be further written as

$$\begin{aligned} & \mathbb{E}_{\hat{p}(\mathbf{x}), p(\mathbf{s}, \kappa | \mathbf{x}; \mathbf{H}^{\text{old}})} \left[\log \frac{p(\mathbf{x}, \mathbf{s}, \kappa; \mathbf{H}^{\text{new}})}{p(\mathbf{s}, \kappa | \mathbf{x}; \mathbf{H}^{\text{old}})} \right] \\ &= \mathcal{Q}(\mathbf{H}^{\text{new}}; \mathbf{H}^{\text{old}}) + \mathbb{E}_{\hat{p}(\mathbf{x}), p(\mathbf{s}, \kappa | \mathbf{x}; \mathbf{H}^{\text{old}})} \left[p(\mathbf{s}, \kappa | \mathbf{x}; \mathbf{H}^{\text{old}}) \right]. \end{aligned}$$

According to subsection 3.2, the optimization problems give

$$\begin{aligned} & \mathcal{Q}(\mathbf{q}^{\text{new}}; \mathbf{H}^{\text{old}}) \geq \mathcal{Q}(\mathbf{q}^{\text{old}}; \mathbf{H}^{\text{old}}), \\ & \mathcal{Q}(\mathbf{A}^{\text{new}}; \mathbf{H}^{\text{old}}) \geq \mathcal{Q}(\mathbf{A}^{\text{old}}; \mathbf{H}^{\text{old}}), \\ & \mathcal{Q}(\mathbf{\Pi}^{\text{new}}; \mathbf{H}^{\text{old}}) \geq \mathcal{Q}(\mathbf{\Pi}^{\text{old}}; \mathbf{H}^{\text{old}}), \\ & \mathcal{Q}(\mathbf{\Theta}^{\text{new}}; \mathbf{H}^{\text{old}}) \geq \mathcal{Q}(\mathbf{\Theta}^{\text{old}}; \mathbf{H}^{\text{old}}). \end{aligned}$$

Since

$$\begin{aligned} \mathcal{Q}(\mathbf{H}^{\text{new}}; \mathbf{H}^{\text{old}}) &= \mathcal{Q}(\mathbf{q}^{\text{new}}; \mathbf{H}^{\text{old}}) + \mathcal{Q}(\mathbf{A}^{\text{new}}; \mathbf{H}^{\text{old}}) \\ &+ \mathcal{Q}(\mathbf{\Pi}^{\text{new}}; \mathbf{H}^{\text{old}}) + \mathcal{Q}(\mathbf{\Theta}^{\text{new}}; \mathbf{H}^{\text{old}}), \end{aligned}$$

it gives

$$\mathcal{Q}(\mathbf{H}^{\text{new}}; \mathbf{H}^{\text{old}}) \geq \mathcal{Q}(\mathbf{H}^{\text{old}}; \mathbf{H}^{\text{old}}).$$

With the above inequality, and the fact that $\mathbb{E}_{\hat{p}(\mathbf{x}), p(\mathbf{s}, \kappa | \mathbf{x}; \mathbf{H}^{\text{old}})} [p(\mathbf{s}, \kappa | \mathbf{x}; \mathbf{H}^{\text{old}})]$ is independent of \mathbf{H}^{new} , we have the inequality

$$\begin{aligned} & \mathbb{E}_{\hat{p}(\mathbf{x}), p(\mathbf{s}, \kappa | \mathbf{x}; \mathbf{H}^{\text{old}})} \left[\log \frac{p(\mathbf{x}, \mathbf{s}, \kappa; \mathbf{H}^{\text{new}})}{p(\mathbf{s}, \kappa | \mathbf{x}; \mathbf{H}^{\text{old}})} \right] \\ & \geq \mathbb{E}_{\hat{p}(\mathbf{x}), p(\mathbf{s}, \kappa | \mathbf{x}; \mathbf{H}^{\text{old}})} \left[\log \frac{p(\mathbf{x}, \mathbf{s}, \kappa; \mathbf{H}^{\text{old}})}{p(\mathbf{s}, \kappa | \mathbf{x}; \mathbf{H}^{\text{old}})} \right]. \end{aligned}$$

Due to $KL(p(\mathbf{s}, \kappa | \mathbf{x}; \mathbf{H}^{\text{old}}) || p(\mathbf{s}, \kappa | \mathbf{x}; \mathbf{H}^{\text{old}})) = 0$, we have

$$\begin{aligned} & \mathbb{E}_{\hat{p}(\mathbf{x})} [\log p(\mathbf{x}; \mathbf{H}^{\text{new}})] \\ & \geq \mathbb{E}_{\hat{p}(\mathbf{x}), p(\mathbf{s}, \kappa | \mathbf{x}; \mathbf{H}^{\text{old}})} \left[\log \frac{p(\mathbf{x}, \mathbf{s}, \kappa; \mathbf{H}^{\text{old}})}{p(\mathbf{s}, \kappa | \mathbf{x}; \mathbf{H}^{\text{old}})} \right] \\ & + \mathbb{E}_{\hat{p}(\mathbf{x})} \left[KL(p(\mathbf{s}, \kappa | \mathbf{x}; \mathbf{H}^{\text{old}}) || p(\mathbf{s}, \kappa | \mathbf{x}; \mathbf{H}^{\text{old}})) \right] \\ & = \mathbb{E}_{\hat{p}(\mathbf{x})} [\log p(\mathbf{x}; \mathbf{H}^{\text{old}})]. \end{aligned}$$

Since $\mathbf{f}_{s,\kappa}$ and $\nabla \mathbf{f}_{s,\kappa}$ are continuous w.r.t. $\theta_{s,\kappa}$ in GenHMM, $\mathbb{E}_{\hat{p}(\mathbf{x})} [\log p(\mathbf{x}; \mathbf{H})]$ is bounded. The above inequality shows $\mathbb{E}_{\hat{p}(\mathbf{x})} [\log p(\mathbf{x}; \mathbf{H})]$ is non-decreasing in learning of GenHMM. Therefore, GenHMM will converge. \square

3.4 Algorithm of GenHMM

Algorithm 1 Learning of GenHMM

- 1: **Input:** Empirical distribution $\hat{p}(\mathbf{x})$ of dataset
- 2: Initializing $\mathbf{H}^{\text{old}}, \mathbf{H} \in \mathcal{H}$ gives:
 $\mathbf{H}^{\text{old}} = \{\mathcal{S}, \mathbf{q}^{\text{old}}, \mathbf{A}^{\text{old}}, p(\mathbf{x} | \mathbf{s}; \Phi_s^{\text{old}})\},$
 $\mathbf{H} = \{\mathcal{S}, \mathbf{q}, \mathbf{A}, p(\mathbf{x} | \mathbf{s}; \Phi_s)\},$
in which generators $\{\mathbf{g}_{s,\kappa} | s \in \mathcal{S}, \kappa = 1, 2, \dots, K\}$ are all initialized randomly.
- 3: $\mathbf{H}^{\text{old}} \leftarrow \mathbf{H}$
- 4: Set learning rate η , neural network optimization batches N per EM step
- 5: **for** \mathbf{H} not converge **do**
- 6: **for** epoch $n < N$ **do**
- 7: Sample a batch of data $\{\mathbf{x}^r\}_{r=1}^{R_b}$ from dataset $\hat{p}(\mathbf{x})$ with batch size R_b
- 8: Compute posterior $p(s_t^r, \kappa_t^r | \mathbf{x}^r; \mathbf{H}^{\text{old}})$
- 9: Formulate loss $\mathcal{Q}(\Theta, \mathbf{H}^{\text{old}})$ in (19)
- 10: $\partial \Theta \leftarrow \nabla_{\Theta} \mathcal{Q}(\Theta, \mathbf{H}^{\text{old}})$
- 11: $\Theta \leftarrow \Theta + \eta \cdot \partial \Theta$
- 12: **end for**
- 13: $\mathbf{q} \leftarrow \underset{\mathbf{q}}{\text{argmax}} \mathcal{Q}(\mathbf{q}; \mathbf{H}^{\text{old}})$ by (28)
- 14: $\mathbf{A} \leftarrow \underset{\mathbf{A}}{\text{argmax}} \mathcal{Q}(\mathbf{A}; \mathbf{H}^{\text{old}})$ by (32)
- 15: $\mathbf{\Pi} \leftarrow \underset{\mathbf{\Pi}}{\text{argmax}} \mathcal{Q}(\mathbf{\Phi}; \mathbf{H}^{\text{old}})$ by (24)
- 16: $\mathbf{H}^{\text{old}} \leftarrow \mathbf{H}$
- 17: **end for**

To summarize the learning solution in subsection 3.2, we wrap our algorithm into pseudocode as shown in Algorithm 1. We use Adam [16] optimizer for optimization w.r.t. the parameters of generators in GenHMM. As shown from line 6 to 10 in Algorithm 1, the batch-size stochastic gradient decent can be naturally embedded into the learning algorithm of GenHMM.

As described by the pseudocode in Algorithm 1, the learning of GenHMM is divided into optimizations w.r.t. to generators' parameters Θ , initial probability q of hidden state, transition matrix \mathbf{A} , and generator mixture weights $\mathbf{\Pi}$. Different from the optimization w.r.t. to q , \mathbf{A} and $\mathbf{\Pi}$, which have optimal solutions, generator learning usually cannot give optimal solution to problem $\max_{\Theta} Q(\Theta; \mathbf{H}^{\text{old}})$. In fact, given that no optimal Θ is obtained, learning of GenHMM can still converge as long as quantity $Q(\Theta; \mathbf{H})$ are improving in iterations in Algorithm 1, where the inequalities in Proposition 1 still hold. Therefore optimal Θ in each iteration is not required for convergence of GenHMM as long as the loss in (19) is getting improved.

4 Experiments

To show the validity of our model, we implement our model in PyTorch and test it with sequential data. We first discuss the experimental setups and then show the experimental results. Code for experiments is available at <https://github.com/FirstHandScientist/genhmm>.

4.1 Experimental Setup

The dataset used for sequential data modeling and classification is TIMIT where the speech signal is sampled at 16kHz. The TIMIT dataset consists of 5300 phoneme-labeled speech utterances which are partitioned into two sets: a train set consists of 4620 utterance, and a test set consists of 1680 utterances. There are totally 61 different types of phones in TIMIT. We performed experiments in two cases: i) full 61-phoneme classification case; ii) 39-phoneme classification case, where 61 phonemes are folded onto 39 phonemes as described in [22].

For extraction of feature vectors, we use 25ms frame length and 10ms frame shift to convert sound track into standard Mel-frequency cepstral coefficients (MFCCs) features. Experiments using the deltas and delta-deltas of the features are also carried out.

Our experiments are performed for: i) standard classification tasks (Table 2, 3, 4, 5), ii) classification under noise perturbation (table 6, 7). The criterion used to report the results includes accuracy, precision and F1 scores. In all experiments, generators $\{g_{s,\kappa} | s \in \mathcal{S}, \kappa = 1, 2, \dots, K\}$ of GenHMM are implemented as flow models. Specifically, our generator structure follows that of a RealNVP described in [8]. As discussed, the coupling layer shown in (11) maps a part of its input signal identically. The implementation is such that layer $l + 1$ would alternate the input signal order of layer l such that no signal remains the same after two consecutive coupling layers. We term such a pair of consecutive coupling layers as a *flow block*. In our experiments, each generator $g_{s,\kappa}$ consists of four *flow blocks*. The density of samples in the latent space is defined as Normal, i.e. $p_{s,\kappa}(z)$ is the density function of standard Gaussian. The configuration for each generator is shown as Table 1.

For each GenHMM, the number of states is adapted to the training dataset. The exact number of states is decided by computing the average length of MFCC frames per phone in training dataset, and clipping the average length into $\{3, 4, 5\}$. Transition matrix \mathbf{A} is initialized as upper triangular matrix for GenHMM.

Table 1: Configuration of generators of GenHMM in Experiments

Latent distribution $p_{s,\kappa}(z)$ $s \in \mathcal{S}, \kappa = 1, 2, \dots, K$	Standard Gaussian
Number of flow blocks	4
Non-linear mapping m_a, m_b	Multiple layer perception, 3 layers and with hidden dimension 24

Table 2: Test accuracy table for 39 dimensional features and folded 39 phonemes.

Model	Criterion	K=1	K=3	K=5
GMM-HMM	Accuracy	62.3%	68.0%	68.7%
	Precision	67.9%	72.6%	73.0%
	F1	63.7%	69.1%	69.7%
GenHMM	Accuracy	76.7%	77.7%	77.7%
	Precision	76.9%	78.1%	78.0%
	F1	76.1%	77.1%	77.0%

Table 3: Test accuracy table for 39 dimensional features and 61 phonemes.

Model	Criterion	K=1	K=3	K=5
GMM-HMM	Accuracy	53.6%	59.6%	61.9%
	Precision	59.1%	63.9%	65.7%
	F1	54.7%	60.5%	62.7%
GenHMM	Accuracy	69.5%	70.6%	70.7%
	Precision	69.2%	70.5%	71.0%
	F1	68.6%	69.6%	69.6%

4.2 Experimental Results

We firstly show the phoneme classification using 39 dimensional MFCC features (MFCC coefficients, deltas, and delta-deltas), to validate one possible usage of our proposed model. Since generative training is carried out in our experiments, GMM-HMM is trained and tested as a reference model in our experiments. Training and testing of GMM-HMM is in the same condition as GenHMMs are trained and tested. Dataset usage for GenHMM and GMM-HMM is the same, and number of states for GMM-HMM is the same as that for GenHMM in modeling each phoneme. Apart from setting the reference model, we also run the experiment comparisons with different total number of mixture components.

Table 2 and 3 shows the results for this experiments, in which we test both the folded 39-phoneme classification case (the conventional way) in Table 2 and the 61-phoneme classification case in Table 3. As shown in both 61-phoneme and 39-phoneme cases, GenHMM gets significant higher accuracy than GMM-HMM for the same number of mixture components. The comparisons with regarding to precision and F1 scores show similar trends and also demonstrate significant improvement of GenHMM's performance. As our expectation, GenHMM has better modeling capacity of sequential data since we bring in the neural network based generators into GenHMM, which should be able to represent complex relationship between states of HMM and sequential data. Apart from the gain of using neural network based generative models, there are also increases of accuracy, precision and F1 scores as the number of mixture components in GenHMM is increased from $K = 1$ to $K = 5$. The sequential dependency of data is modeled by HMM itself, while each state of HMM can have better representation using a mixture probabilistic model if data represented by the state is multi-mode. Comparing the results in 39-phoneme and 61-phoneme cases, GenHMM gets higher accuracy

for 39-phoneme classification than it does for 61-phoneme classification. The total training dataset size remains the same as 61 phonemes are folded into 39 phonemes. There are less training data available per phonemes and more classes to be recognized in the 61-phoneme case, which makes the task more challenging.

Table 4: Test accuracy table for 13 dimensional features and folded 39 phonemes.

Model	Criterion	K=1	K=3	K=5
GMM-HMM	Accuracy	48.5%	51.2%	52.4%
	Precision	56.2%	58.3%	59.5%
	F1	50.3%	53.0%	54.2%
GenHMM	Accuracy	61.1%	62.1%	62.1%
	Precision	61.1%	61.9%	62.1%
	F1	59.7%	60.7%	60.2%

Table 5: Test accuracy table for 13 dimensional features and 61 phonemes.

Model	Criterion	K=1	K=3	K=5
GMM-HMM	Accuracy	37.1%	40.6%	42.2%
	Precision	44.6%	47.4%	48.8%
	F1	38.8%	42.1%	43.7%
GenHMM	Accuracy	50.3%	50.8%	52.3%
	Precision	49.3%	50.9%	52.1%
	F1	47.8%	48.3%	49.3%

Similar experiments are carried out by using only the MFCC coefficients as feature input (excluding deltas and delta-deltas). The results are shown in Table 4 and 5. The superior performance of GenHMM remains compared with reference model GMM-HMM, with regarding to accuracy, precision and F1 scores. The gain by using mixture generators is also presented in this set of experiments while the difference between 61-phoneme and 39-phoneme cases is similar to the set of experiments in Table 2 and 3.

Table 6: Test accuracy table of perturbation with white noise ($K = 3$, folded 39 phonemes).

Model	Criterion	White Noise SNR			
		15dB	20dB	25dB	30dB
GMM-HMM	Accuracy	36.6%	44.2%	50.8%	57.1%
	Precision	59.2%	64.2%	68.4%	70.6%
	F1	39.9%	47.7%	53.9%	59.9%
GenHMM	Accuracy	52.4%	62.0%	69.7%	74.3%
	Precision	60.0%	65.9%	71.7%	74.8%
	F1	52.5%	62.0%	69.3%	73.5%

Apart from standard classification testing, we also test the robustness of our model to noise perturbations. We train GenHMM with $K = 3$ by clean TIMIT training data in the case of folded 39 phonemes with 39 dimensional features. The testing dataset is perturbed by either the same type of noise with different signal-to-noise ratio (SNR) as shown in Table 6, or different type of noises with the same SNR as shown in Table 7. The noise data is from NOISEX-92 database. The baseline of these two sets of experiments is the accuracy testing of GenHMM and GMM-HMM on clean testing data in the same experimental condition, where GenHMM has 77.7% and

Table 7: Test accuracy table of perturbation by different type of noise (SNR=20dB, $K = 3$, folded 39 phonemes).

Model	Criterion	Noise Type			
		White	Pink	Babble	Volvo
GMM-HMM	Accuracy	44.2%	48.8%	57.7%	66.6%
	Precision	64.2%	66.1%	67.0%	71.9%
	F1	47.7%	52.3%	59.7%	67.8%
GenHMM	Accuracy	62.0%	65.1%	70.0%	75.7%
	Precision	65.9%	67.8%	70.4%	75.9%
	F1	62.0%	64.6%	69.0%	75.3%

GMM-HMM gets 68.0% as shown in Table 2. Similar superior performance of GenHMM with regarding to precision and F1 scores is also shown. It is shown in Table 6 that GMM-HMM’s performance degenerates more than GenHMM’s performance at the same level of noise perturbation, though the accuracy of both models increases along the increase of SNR. Especially, for SNR=30dB, the accuracy of GenHMM drops only about 3% (from 77.7% to 74.3%), while GMM-HMM encounters more than 10% decrease (from 68.0% to 57.1%) due to the noise perturbation. In Table 7, the SNR remains constant and GenHMM is tested with perturbation of different noise types. It is shown that GenHMM still remain higher performance scores at different types of noise perturbations than GMM-HMM. Among these four types of noise, white noise shows most significant impact to GenHMM while the impact of volvo noise is negligible.

5 Conclusion

In this work, we proposed a generative model based HMM (GenHMM) whose generators are realized by neural networks. We provided the training method for GenHMM. The validity of GenHMM was demonstrated by the experiments of classification tasks on practical sequential dataset. The learning method in this work is based on generative training. For future work, we would consider discriminative training for classification tasks of sequential data.

ACKNOWLEDGEMENTS

We would like to thank Dr. Minh Thành Vu for his discussions and comments on the algorithm analysis, which helped improve this paper considerably.

REFERENCES

- [1] N.M.R. Ashwin, Leonard Barnabas, Amalraj Ramesh Sundar, Palaniyandi Malathi, Rasappa Viswanathan, Antonio Masi, Ganesh Kumar Agrawal, and Randeep Rakwal, ‘Comparative secretome analysis of colletotrichum falcatum identifies a cerato-platanin protein (ep11) as a potential pathogen-associated molecular pattern (pamp) inducing systemic resistance in sugarcane’, *Journal of Proteomics*, **169**, 2 – 20, (2017). 2nd World Congress of the International Plant Proteomics Organization.
- [2] Bing-Hwang Juang, S. Levinson, and M. Sondhi, ‘Maximum likelihood estimation for multivariate mixture observations of markov chains (corresp.)’, *IEEE Transactions on Information Theory*, **32**(2), 307–309, (March 1986).
- [3] Christopher M. Bishop, *Pattern Recognition and Machine Learning (Information Science and Statistics)*, Springer-Verlag, Berlin, Heidelberg, 2006.
- [4] Jan Buys, Yonatan Bisk, and Yejin Choi, ‘Bridging hmms and rnns through architectural transformations’, in *32nd Conference on Neural Information Processing Systems, IRASL workshop*, (2018).

- [5] S. Chatterjee and W. B. Kleijn, 'Auditory model-based design and optimization of feature vectors for automatic speech recognition', *IEEE Transactions on Audio, Speech, and Language Processing*, **19**(6), 1813–1825, (Aug 2011).
- [6] Tian Qi Chen, Yulia Rubanova, Jesse Bettencourt, and David K Duvenaud, 'Neural ordinary differential equations', in *Advances in Neural Information Processing Systems 31*, eds., S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, 6571–6583, Curran Associates, Inc., (2018).
- [7] W. Ding, S. Li, H. Qian, and Y. Chen, 'Hierarchical reinforcement learning framework towards multi-agent navigation', in *2018 IEEE International Conference on Robotics and Biomimetics (ROBIO)*, pp. 237–242, (Dec 2018).
- [8] L. Dinh, J. Sohl-Dickstein, and S. Bengio, 'Density estimation using Real NVP', *ArXiv e-prints*, (May 2016).
- [9] Laurent Dinh, David Krueger, and Yoshua Bengio, 'NICE: non-linear independent components estimation', *CoRR*, **abs/1410.8516**, (2014).
- [10] M. Gales and S. Young, *Application of Hidden Markov Models in Speech Recognition*, now, 2008.
- [11] Will Grathwohl, Ricky T. Q. Chen, Jesse Bettencourt, Ilya Sutskever, and David Duvenaud, 'FFJORD: free-form continuous dynamics for scalable reversible generative models', *CoRR*, **abs/1810.01367**, (2018).
- [12] Trienani Hariyanti, Saori Aida, and Hiroyuki Kameda, 'Samawa language part of speech tagging with probabilistic approach: Comparison of unigram, HMM and TnT models', *Journal of Physics: Conference Series*, **1235**, 012013, (jun 2019).
- [13] G. Hinton, L. Deng, D. Yu, G. E. Dahl, A. Mohamed, N. Jaitly, A. Senior, V. Vanhoucke, P. Nguyen, T. N. Sainath, and B. Kingsbury, 'Deep neural networks for acoustic modeling in speech recognition: The shared views of four research groups', *IEEE Signal Processing Magazine*, **29**(6), 82–97, (Nov 2012).
- [14] Geoffrey E. Hinton, *A Practical Guide to Training Restricted Boltzmann Machines*, 599–619, Springer Berlin Heidelberg, Berlin, Heidelberg, 2012.
- [15] Wahab Khan, Ali Daud, Jamal A Nasir, and Tehmina Amjad, 'A survey on the state-of-the-art machine learning models in the context of nlp', *Kuwait journal of Science*, **43**(4), (2016).
- [16] Diederik P. Kingma and Jimmy Ba, 'Adam: A method for stochastic optimization', *CoRR*, **abs/1412.6980**, (2014).
- [17] Durk P Kingma and Prafulla Dhariwal, 'Glow: Generative flow with invertible 1x1 convolutions', in *Advances in Neural Information Processing Systems 31*, eds., S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, 10215–10224, Curran Associates, Inc., (2018).
- [18] Viktoriya Krakovna and Finale Doshi-Velez, 'Increasing the Interpretability of Recurrent Neural Networks Using Hidden Markov Models', *arXiv e-prints*, arXiv:1606.05320, (Jun 2016).
- [19] Sergey Levine, 'Reinforcement learning and control as probabilistic inference: Tutorial and review', *CoRR*, **abs/1805.00909**, (2018).
- [20] L. Li, Y. Zhao, D. Jiang, Y. Zhang, F. Wang, I. Gonzalez, E. Valentin, and H. Sahli, 'Hybrid deep neural network–hidden markov model (dnn-hmm) based speech emotion recognition', in *2013 Humaine Association Conference on Affective Computing and Intelligent Interaction*, pp. 312–317, (Sep. 2013).
- [21] Larkin Liu, Yu-Chung Lin, and Joshua Reid, 'Improving the performance of the LSTM and HMM models via hybridization', *CoRR*, **abs/1907.04670**, (2019).
- [22] Carla Lopes and Fernando Perdigao, 'Phoneme recognition on the timit database', in *Speech Technologies*, ed., Ivo Ipsic, chapter 14, IntechOpen, Rijeka, (2011).
- [23] Yajie Miao and Florian Metze, 'Improving low-resource cd-dnn-hmm using dropout and multilingual dnn training', in *INTERSPEECH*, (2013).
- [24] S. Ren, V. Sima, and Z. Al-Ars, 'Fpga acceleration of the pair-hmms forward algorithm for dna sequence analysis', in *2015 IEEE International Conference on Bioinformatics and Biomedicine (BIBM)*, pp. 1465–1470, (Nov 2015).